

James Sethian Receives Norbert Wiener Prize in Applied Mathematics

James Sethian, head of the Mathematics Group in CRD and a professor of mathematics at the University of California, Berkeley, has been awarded the Norbert Wiener Prize in Applied Mathematics by the American Mathematical Society (AMS) and the Society for Industrial and Applied Mathematics (SIAM). The prize, presented Jan. 8 at the joint AMS-SIAM meeting in Phoenix, is awarded for an outstanding contribution to "applied mathematics in the highest and broadest sense."

Sethian's award marks the eighth time the Norbert Wiener Prize has been awarded since 1970. The prize was last awarded in 2000, and one of the two recipients was Alexandre Chorin, a colleague of Sethian's who also has a joint appointment in LBNL's Mathematics Group and UC Berkeley's Math Department.

According to information distributed at the AMS-SIAM meeting awards ceremony, Sethian was honored "for his seminal work on the computer representation of the motion of curves, surfaces, interfaces, and wave fronts, and for his brilliant applications of mathematical and computational ideas to problems in science and engineering."

His work has influenced fields as diverse as medical imaging, seismic research by the petroleum industry, and the manufacture of computer chips and desktop printers. AMS and SIAM provided the following descriptions of Sethian's work and its importance:

(continued on page 3)


Introducing CRD Report

CRD Report is a new publication highlighting recent achievements by staff members in Berkeley Lab's Computational Research Division. Initially, CRD Report will be distributed every other month via email and posted on the Web at <http://crd.lbl.gov/DOEResources>. It may be freely distributed. CRD Report is edited by Jon Bashor, JBashor@lbl.gov or 510-486-5849.

Improved Algorithms Lead to Lab-Scale Combustion Simulations

In spite of its fundamental technological importance, our knowledge of basic combustion processes is surprisingly incomplete. Theoretical combustion science is unable to address the complexity of realistic flames, and laboratory measurements are difficult to interpret and often limited in the types of applicable flames or levels of detail they can provide. CRD's Center for Computational Sciences and Engineering (CCSE) has teamed with LBNL's Environmental Energy Technologies Division (EETD) to build a high-performance computing solution to flame simulation and analysis that has a unique potential for making dramatic progress in combustion science research.

(continued on page 3)



DSD Tools Making Grids Easier to Use

At the SC2003 conference held last November in Phoenix, Grids were omnipresent, with speakers and vendors emphasizing the development and deployment of Grid infrastructure and applications. For all their potential, though, actually using Grids can be difficult and time-consuming.

That's where Berkeley Lab's Distributed Systems Department (DSD) (<http://dsd.lbl.gov>) comes in. At the conference, DSD staff members demonstrated "wxPyGlobusJobGui", a graphical user interface which integrates functionality from three DSD projects to provide a securely authenticated, Grid-enabled and network-monitored prototype job submission and monitoring system. WxPyGlobusJobGui demonstrated the utility of pyGlobus to rapidly generate a graphical user interface which automated the steps necessary to stage (copy) input and

control files from the SC conference exhibition floor to a cluster located at Berkeley Lab, run the AMBER molecular dynamics application on the input files, and stage the result files back to SC03. The graphical user interface demonstrated two other DSD-developed technologies: Akenti and NetLogger.

Akenti (<http://dsd.lbl.gov/akenti/>) is an access control policy library which uses digitally signed certificates to define access policies for shared resources. Integrating the Akenti access control policy library into the Globus Gatekeeper and Jobmanager provided an access control mechanism significantly more powerful than the local user/certificate mappings currently used.

NetLogger (<http://dsd.lbl.gov/NetLogger>) is a lightweight toolkit which facilitates the analysis and debugging of distributed computing environments. Adding NetLogger logging calls to

(continued on page 2)

NEWS

CRD's Brian Tierney Uses Expertise to Improve Performance of Networked Systems

A good mechanic doesn't just keep a car running, but has a toolkit of expertise and expert tricks to get the most performance out of a vehicle. And so it is with networks and computers, especially when distributed computing sources



Brian Tierney

are linked by networks to perform as an integrated system. Brian Tierney of the Distributed Systems Department has built a reputation as an expert in finding new ways to improve the performance of distributed systems. He's shared his talents with DARPA, Internet2, CERN

and other national labs.

Next up on his agenda is PFDLnet, the Second International Workshop on Protocols for Long-Distance Networks (<http://www-didc.lbl.gov/PFDLnet2004/index.htm>), which Brian is co-chairing with Les Cottrell of SLAC

on February 16-17. At the heart of the workshop is a known problem and four potential solutions.

The problem is with TCP, the transmission control protocol that ensures that data packets are delivered as sent. While TCP is one of the key elements of the Internet's initial success, it's now become a speedbump. TCP just doesn't scale well to work on really fast, long-distance networks – high bandwidth and high latency overwhelm the congestion control algorithm in the protocol. A number of techniques have been developed to deal with the problem, and Tierney himself has created a Web page with 100 tips and tricks for tuning TCP to improve performance (go to <http://www-didc.lbl.gov/TCP-tuning/TCP-tuning.html>).

"When trying to do Grid troubleshooting, one of the first things people do is blame the network – and that's true less than half the time," Tierney said. "You can tune the network in terms of TCP, but we've exhausted the limits of tuning. For very high-speed networks, a new protocol is really needed."

The goal of the workshop is to try to come up with a consensus on how to proceed. On the agenda are four approaches, ranging from minor tweaks of the current TCP to a complete overhaul. The tactics range from enhancing the current TCP congestion response algorithms, congestion avoidance-based techniques, building reliability on top of UDP (User Datagram Protocol, or the Unreliable Data Protocol, depending on your perspective) to radically new protocols such as XCP (eXplicit control protocol) which requires router support so that the routers inform the sender of the level of congestion at the bottle neck, allowing the sender to back off before congestion (loss) occurs.

While 18 papers will be presented over the two days of the workshop, there will also be plenty of time for discussion, which Tierney said is critical to the success of the workshop.

Tierney stepped forward to co-chair this workshop after attending the first one, and he decided that it was important for the community. "It's very appropriate that DOE take a leadership role in this work as we have the big applications, and the Science Grid, that TCP is now getting in the way of," he said.

Contact Brian at BLTierney@lbl.gov.

DSD Tools (continued from p.1)


all the elements in the demonstration, from the file staging and job submission components to the Globus Gatekeeper/Jobmanager, the Akenti authorization code, and the scientific application, streamlined the process of tracking down problems that occurred during the demonstration, such as invalid input file parameters, authorization failures, and transient problems. NetLogger was quite effective in helping trace concurrent local and remote operations.

By using existing resources such as the Python scripting language, the pyGlobus toolkit (<http://dsd.lbl.gov/gtg/projects/pyGlobus/>), the NetLogger toolkit, and the wxWindows GUI library, DSD staff members were able to develop the user interface and functional components of the demonstration application extremely rapidly – in less than two weeks. Also, adding NetLogger messages to the pyGlobus file transfer and job submission modules was significantly easier than adding them to the source code of the underlying Globus toolkit.

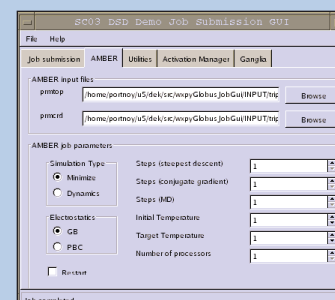
While demonstrating this work at the conference, staff members learned that there is significant interest in having generic toolkits which can be specialized. Although this demo focused on AMBER, the demo GUI was designed to allow other domain-specific applications to re-use the file staging and job submission functionality.

Already, ideas for improving the application are being discussed. It should be possible to describe significantly more sophisticated and dynamic data and computational operations. Many computational procedures involve the same elements, such as file staging and remote job execution; chains of such elements, constructed to perform common scientific jobs, are called "workflows." The application would also be enhanced with a mechanism in which these workflows can be visually described, launched, and monitored.

For more information, contact Keith Jackson, krjackson@lbl.gov or go to <http://dsd.lbl.gov/sc03/>



The SC2003 DSD Job Submission timeline GUI reflecting Netlogger "events"



The application-specific properties page for the AMBER molecular dynamics application.

Sethian Wins Prize (continued from p.1)


"A particularly noteworthy aspect of Sethian's work is that he pursues his ideas from a first formulation of a mathematical model all the way to concrete applications in national laboratory and industrial settings; his algorithms are currently distributed in widely available packages," the AMS and SIAM noted.

"Sethian's work is a shining example of what applied mathematics can accomplish to benefit science as a whole.

"Sethian's earliest work included an analysis of the motion of flame fronts and of the singularities they develop; he found important new links between the motion of the fronts and partial differential equations. These con-

nections made possible the development of advanced methods to describe front propagation through the solution of regularized equations on fixed grids.

"Sethian (with S. Osher) extended this work through an implicit formulation, resulting in a methodology that has come to be known as the 'level set method,' because it represents a front propagating in n dimensions as a level set of an object in $(n+1)$ dimensions. Next, Sethian tamed the cost of working in higher dimensions by reducing the problem back down to its original dimensionality. This set of ideas makes possible the solution of practical problems of increasing importance and sophistication and constitutes a major mathematical development as well as an



exceptionally useful computational tool with numerous applications.

Sethian is also the author of a book entitled "Level Set Methods and Fast Marching Methods" published by Cambridge University Press.


Contact: JASethian@lbl.gov. More information about Sethian's work can be found at <http://math.berkeley.edu/~sethian/>

Combustion Research

(continued from p.1)

The CCSE group has created the first detailed simulations of laboratory-scale turbulent premixed flame experiments using 3D time-dependent software accommodating an unprecedented level of detail in terms of chemical fidelity and fluid transport. CCSE and EETD researchers work together to validate the simulations with experimental data, and then probe the computed results for information not easily obtainable from the experiment in any other way. The investigations are focused in two primary areas: how turbulence in the fuel stream affects the local combustion chemistry, and how emissions are formed and released in the product stream. The work has applications for devices ranging from power generators to heating systems, water heaters, stoves, ovens and even clothes driers.

Simulation at such a level of detail was impossible just a few years ago. However, algorithmic improvements by DOE-funded applied mathematics groups such as CCSE over the past five years have slashed computational costs for these types of flows by a factor of 10,000. Such savings enable key improvements in the fidelity of the chemical and fluid dynamical descriptions of the flows, to the point that real experiments may now be simulated without ad-hoc engineering models for under-resolved physical processes. However, simulation of practical-scale combustion devices remains an immense undertaking. CCSE has implemented their advanced simulation algorithms on state-of-the-art parallel computing hardware to



These methane flame images show exceptional agreement between CCSE's simulation (left) and experimental data (right).

increase the number of variables available for describing the system from hundreds of thousands five years ago, to more than a billion today.

The research approach taken by CCSE has explicitly targeted both the temporal and spatial multi-scale aspects of combustion modeling. The group takes advantage of key mathematical characteristics of low-speed flows, common to most combustion applications, to eliminate components of the model relevant only to high-speed scenarios. For low-speed flows, these components have little effect on the system dynamics, yet they drive down the simulation efficiency by unnecessarily limiting the maximum numerical time-step size. The integration algorithms are implemented in a set of software tools based on adaptive mesh refinement (AMR), a dynamic grid-based system that automatically allocates computational resources to regions that contain the most interesting details. The AMR methodology allows one to simultaneously incorporate large-scale effects that stabilize the flame, as well as the very fine-scale features of the combustion reaction zone itself.

The detailed solutions computed by CCSE are being validated with experimental data provided by the EETD Combustion Lab. Comparisons include global observables, such as mean flame locations and geometries, as well as statistics of instantaneous flame surface structures. In addition to simply validating the computed solutions, however, the research groups probe the massive amounts of data generated by the computation in order to learn more about flame details, such as the localized effects of large and small eddies on the structure of the combustion reaction zone. For example, the distribution of hydrogen atoms in the thermal field is tightly coupled to key chain-branching reactions required to sustain the combustion process itself. The detailed models accurately represent the transport of hydrogen with respect to the other chemical species in the context of this turbulent flow. CCSE is presently using detailed chemistry and transport models containing 20 to 65 chemical species and hundreds of reactions.

CCSE is now working with EETD researchers to develop statistical measures of the simulation and experimental data so that they can obtain a more quantitative comparison. They are also working to understand the volumes of new simulation data in order to quantitatively characterize how fuel-stream turbulence affects the detailed combustion process.

For more information about the CCSE's adaptive methodology for low-Mach number combustion modeling applications, visit <http://seesar.lbl.gov/CCSE> or contact Marc Day at MSDay@lbl.gov or John Bell at JBBell@lbl.gov.